arXiv:1703.09700v1 [cs.LG] 28 Mar 2017

Inverse Reinforcement Learning from Summary Data

Antti Kangasrääsiö [1] **Samuel Kaski** [1]

Abstract

Inverse reinforcement learning (IRL) aims to explain observed complex behavior by fitting reinforcement learning models to behavioral data. However, traditional IRL methods are only applicable when the observations are in the form of state-action paths. This is a problem in many real-world modelling settings, where only more limited observations are easily available. To address this issue, we extend the traditional IRL problem formulation. We call this new formulation the inverse reinforcement learning from summary data (IRL-SD) problem, where instead of state-action paths, only summaries of the paths are observed. We propose exact and approximate methods for both maximum likelihood and full posterior estimation for IRL-SD problems. Through case studies we compare these methods, demonstrating that the approximate methods can be used to solve moderate-sized IRL-SD problems in reasonable time.

1. Introduction

Inverse reinforcement learning (IRL) has generally been formulated (Russell, 1998; Ng & Russell, 2000; Ramachandran & Amir, 2007; Ziebart et al., 2008) as:

Given (1) a set of state-action paths $\Xi = \{\xi_1, \ldots, \xi_N\}$, where $\xi_i = (s_0^i, a_1^i, \ldots, a_{T_i-1}^i, s_{T_i}^i)$; (2) a Markov decision-process (MDP) with reward-function $R(s, a; \theta)$, where the θ are unknown parameters; and optionally (3) prior $P(\theta)$.[1]
Determine a point estimate $\hat{\theta}$ or a posterior $P(\theta|\Xi)$.

Methods for solving the IRL problem have been used for parameter inference in multiple real-world modelling situations where complex behavior has been observed in the form of state-action paths. Examples include driver route modelling (Ziebart et al., 2008), helicopter acrobatics (Abbeel et al., 2010), learning to perform motor tasks (Boularias et al., 2011), dialogue systems (Chandramohan et al., 2011), pedestrian activity prediction (Ziebart et al., 2009; Kitani et al., 2012), and commuting routines (Banovic et al., 2016). An overview of existing methods is given in Section 2.

However, a limitation with the traditional problem formulation is the assumption that full paths containing both actions and states have been observed.[2] In many real-world situations such fine-grained observations may not be available, or acquiring them may be prohibitively expensive. In fact, it can be argued that the number of cases where only more limited observations are available is probably much larger than those with accurate path data.

It is a common scenario that various other types of observations of the behavior of the agent are available, and we would like to estimate the preferences, tasks or properties of the agent. One such example would be identifying the capabilities of computer users based on observations from them "in the wild", where the types of observations may vary case-by-case (Wobbrock et al., 2011).

To extend IRL to the kind of situations where instead of state-action paths, arbitrary *summaries* of the underlying true paths are observed, we formulate the inverse reinforcement learning from summary data (IRL-SD) problem, defined in Section 3.

In Section 4 we derive the likelihood function for this problem, but observe that evaluating it is computationally expensive. To address this issue, we propose a surrogate that is faster to evaluate in practice. For inference, we propose a Bayesian optimization based method.

In Section 5 we demonstrate that by using the surrogate we get results comparative to using the exact likelihood. We also demonstrate that we are able to infer ML estimates for moderate-sized toy models and the full posterior for a model of human visual search, based on only summary observations.

[1]Helsinki Institute for Information Technology HIIT, Department of Computer Science, Aalto University, Finland. Correspondence to: Antti Kangasrääsiö <antti.kangasraasio@aalto.fi>, Samuel Kaski <samuel.kaski@aalto.fi>.

Copyright 2017 by the author(s).

[1] The notation is elaborated in Section 2.

[2] Or in some formulations, a sufficient amount of separate state-action transitions.

2. Inverse Reinforcement Learning

Here we give a short overview of the general IRL modelling assumptions, and standard assumptions existing methods make of the observation data. We also review the main approaches used for solving the inference problem. A recent review was also made by Gao et al. (2012).

2.1. IRL Model Assumption

The standard IRL modeling assumption is that an agent is interacting with a MDP environment, demonstrating optimal behavior over N independent episodes, thus creating paths $\Xi = (\xi_1, \ldots, \xi_N)$. Each path is a sequence of states and actions, denoted as $\xi_i = (s_0^i, a_0^i, \ldots, a_{T_i-1}^i, s_{T_i}^i)$, where s_t and a_t are the state and action at timestep t, and T_i is the length of trajectory i. We will later omit sub- and superscripts when possible.

MDP M is defined by the tuple (S, A, T, R, γ), where S is a set of states, A is a set of actions, $T = P(s_{t+1}|s_t, a_t)$ is the state transition function, $R(s)$ is the reward function, and γ is the temporal discount rate. Here we assume M to be defined in terms of some unknown parameters θ. An instance of M with fixed parameters θ is denoted by M_θ.

A generative model for simulating observation datasets, denoted $\texttt{RLSIM}(M_\theta, N)$, is illustrated in Algorithm 1. Here we denote by $\texttt{RL}(M_\theta)$ a generic subroutine for solving the optimal policy π^* for M_θ, and by $\texttt{SIM}(M_\theta, \pi)$ a subroutine for simulating paths ξ given a policy π.

Algorithm 1 IRL generative model: RLSIM

Input: M_θ, N
Output: Ξ
$\pi^* \leftarrow \texttt{RL}(M_\theta)$
for $i = 1$ **to** N **do**
  $\xi_i \leftarrow \texttt{SIM}(M_\theta, \pi^*)$
end for
$\Xi \leftarrow (\xi_1, \ldots, \xi_N)$

2.2. Observation Assumptions

There are four different situations that have been studied in more depth in past IRL literature.

(1) The policy π of the agent is known (Ng & Russell, 2000). In other words, we assume that we know exactly how the agent behaves in any situation, but don't know why.

(2) Noise-free observations of the states of the environment and actions of the agent are available (Ng & Russell, 2000; Ratliff et al., 2006; Neu & Szepesvári, 2007; Ramachandran & Amir, 2007; Dimitrakakis & Rothkopf, 2011; Rothkopf & Dimitrakakis, 2011; Klein et al., 2012; Michini & How, 2012; Klein et al., 2013; Tossou & Dimitrakakis, 2013; Choi & Kim, 2015; Nguyen et al., 2015; Herman et al., 2016). This is probably the most common formulation in the literature. A benefit of this assumption is that it allows the likelihood to be factorized per state transition.

(3) Feature expectations of paths traveled by the agent are available (Abbeel & Ng, 2004; Ziebart et al., 2008; Boularias et al., 2011; Bloem & Bambos, 2014). Here feature expectations are defined as $\hat{\mu}_E = \frac{1}{N}\sum_{i=1}^{N}\sum_{t=0}^{T_i}\gamma^t\phi(s_t^i)$, where the underlying assumption is that N state sequences $(s_0^i, \ldots, s_{T_i}^i)$ have taken place, γ is the discount rate of the agent and ϕ is a state feature vector function. This follows from the assumption that the reward function is linear in state features, $R(s) = \theta^T\phi(s)$, after which the inference problem is formulated as a function of $\theta^T\hat{\mu}_E$. Feature expectations can be seen as specific kinds of summaries of the actual paths.

(4) Noisy observations of the states of the environment are available (Kitani et al., 2012; Surana, 2014). Here it is assumed that state sequences $(s_0^i, \ldots, s_{T_i}^i)$ have taken place, but only observations $(u_0^i, \ldots, u_{T_i}^i)$ are available, where u is a probability distribution over the state space, defining the observer's uncertainty of the actual state. The general approach is to use a forward-backward algorithm for estimating the state visitation frequencies and use them in turn to compute the feature expectations $\hat{\mu}_E$. Noisy observations can also be seen as specific kinds of summaries of the actual paths.

2.3. Solution Approaches

There are two common approaches for solving the IRL problem. MCMC can be applied for computing samples of the posterior when the unnormalized likelihood can be evaluated in closed form. A characteristic example of this is the Bayesian IRL method proposed by Ramachandran and Amir (Ramachandran & Amir, 2007).[3] Gradient descent can be applied for giving point estimates when the gradient of the likelihood can be evaluated in closed form. A characteristic example of this is the Maximum Entropy IRL method proposed by Ziebart et al. (Ziebart et al., 2008). Also point estimation based on linear programming (e.g. Ng & Russell, 2000) and classification (e.g. Klein et al., 2012) have been considered.

3. IRL from Summary Data

Here we give our precise problem definition and explain the motivation behind it.

[3] Although MAP estimation can also be performed (Choi & Kim, 2011).

3.1. Problem Definition

Let M be a MDP parametrized by θ, where θ is any finite set of parameters of interest (not limited to the reward function parameters). Let the true parameters be θ^* and assume an agent whose behavior agrees with an optimal policy for M_{θ^*}. We do not know θ^*, but know that the true values reside within a bounded region Θ. We may additionally have a prior $P(\theta)$ that defines how likely each θ is *a priori*.

Assume that the agent has taken paths $(\xi_1, \ldots, \xi_N)$ but we only have observed *summaries* of these paths: $\Xi_\sigma = (\xi_{1\sigma}, \ldots, \xi_{N\sigma})$, where $\xi_{i\sigma} = \sigma(\xi_i)$ and σ is a known summary function, mapping a path into another type of observation[4]. The inverse reinforcement learning problem from summary data (IRL-SD) is then defined as:

Given (1) set of summaries Ξ_σ of an agent demonstrating optimal behavior; (2) summary function σ; (3) MDP M with θ unknown; (4) bounded space Θ; and optionally (5) prior $P(\theta)$.
Estimate either $\hat{\theta} \in \Theta$ such that simulated optimal behavior from $M_{\hat{\theta}}$ agrees with Ξ_σ, or the posterior $P(\theta|\Xi_\sigma)$.

In the traditional IRL setting, θ would be the parameters of the reward function. Our formulation extends the inference problem to any other parameters as well, such as those of the state space, action space, transition function and discount rate.[5] Also, an additional benefit is that the formulation is not limited to just linear reward functions, unlike in most existing methods.

3.2. Motivation

There are multiple situations where we could hope to estimate the parameter values of a RL-based model, but only have convenient access to summary-type observations.

(1) Observing paths is expensive. There are multiple situations where coarse-grained or noisy observations are significantly cheaper to acquire compared to accurate path observations. For example, in human-computer interaction, setting up a user study where eye-tracking and thinking aloud is used to get accurate estimates of the cognitive states and actions of the user is much more expensive compared to logging user's mouse and keyboard activity.

(2) Adversarial situations cause data censoring and distortion. If we are trying to infer the strategy of an opponent, it is likely that she will attempt to hide the details of her states and actions if possible. This is common in games of incomplete information, such as card games (e.g. poker) or real-time strategy games (e.g. Starcraft).

[4] If σ is identity, then this problem definition essentially matches the traditional IRL problem.

[5] Recently Herman et al. (2016) considered the simultaneous inference of rewards and transition dynamics as well.

(3) Ensuring privacy prevents accurate observations. For example, in differentially private situations, our access to the observation data may be limited. We might only have access to summary statistics of the observations, in order to protect the privacy of the data sources.

We also note that the IRL-SD problem is analogous to many natural learning situations. For example, when humans are instructed by a teacher, they might not be able to observe the states and actions experienced by the teacher. Instead the learner might receive, for example, verbal instruction that summarizes the desired behavior. In these types of situations, the human learner might simulate various likely scenarios either physically or mentally, in order to be able to understand the concept being taught.

4. Inference Methods for IRL-SD

In this section we first derive the observation likelihood for the IRL-SD problem. However, as evaluating the likelihood function can be very expensive, we also propose a surrogate that is generally much faster to evaluate. After that, we propose a Bayesian optimization based method for performing inference.

4.1. Exact Likelihood

To derive a computable likelihood, we need to make some assumptions. First, we assume both $|S|$ and $|A|$ are finite (e.g. through discretization). We also assume that the maximum number of actions that can be performed within an observed episode is T_{max}. We denote the finite set of all plausible trajectories (that have non-zero contribution to the likelihood) by $\Xi_{ap} \subseteq S^{T_{max}+1} \times A^{T_{max}}$.

Now the likelihood for θ given summary observations $\Xi_\sigma = (\xi_{1\sigma}, \ldots, \xi_{N\sigma})$ can be written as

$$L(\theta|\Xi_\sigma) = \prod_{i=1}^{N} \big[P(\xi_{i\sigma}|\theta)\big] \tag{1}$$

$$= \prod_{i=1}^{N} \Big[\sum_{\xi_i \in \Xi_{ap}} \big[P(\xi_{i\sigma}|\xi_i) P(\xi_i|\theta)\big] \Big], \tag{2}$$

where

$$P(\xi_{i\sigma}|\xi_i) = P\big(\sigma(\xi_i) = \xi_{i\sigma}\big), \tag{3}$$

and

$$P(\xi_i|\theta) = P(s_0^i) \prod_{t=0}^{T_i - 1} \big[\pi_\theta^*(s_t^i, a_t^i) P(s_{t+1}^i|s_t^i, a_t^i)\big]. \tag{4}$$

The main difficulty with the exact likelihood is finding the set Ξ_{ap} and evaluating the sum over it.

4.2. Approximate Likelihood

For deriving the approximate method, we start with the goal of IRL-SD: "estimate $\hat{\theta} \in \Theta$ such that simulated optimal behavior from $M_{\hat{\theta}}$ agrees with Ξ_σ".

First, in order to validate that we have reached the goal of IRL-SD, we need to be able to simulate summary datasets Ξ_σ^{sim} from optimal behavior, with parameters θ and using the summary function σ. For this purpose, we define $\texttt{RLSUM}(M_\theta, N, \sigma)$ as a modification of $\texttt{RLSIM}$ (Algorithm 1) with the added summarization. We also note that $\texttt{RLSUM}(M_\theta, |\Xi_\sigma|, \sigma)$ is the generative process for Ξ_σ.

Second, we need to define what is meant by two summary datasets "agreeing with" each other. For this purpose, we assume a *discrepancy function* δ exists, such that

$$\delta(\Xi_\sigma^A, \Xi_\sigma^B) \to [0, \infty), \tag{5}$$

where lower values of δ indicate higher level of agreement between the two datasets.

By combining $\texttt{RLSUM}(M_\theta, |\Xi_\sigma|, \sigma)$ with δ, we can define a stochastic variable

$$d_\theta \sim \delta(\texttt{RLSUM}(M_\theta, |\Xi_\sigma|, \sigma), \Xi_\sigma), \tag{6}$$

whose distribution corresponds with the ability of θ to satisfy our requirements for solving the IRL-SD problem we have at hand.

Now we take inspiration from approximate Bayesian computation (ABC) (see Sunnåker et al., 2013), and define our surrogate likelihood function so that it corresponds with the probability of d_θ being "small", according to

$$\tilde{L}_\varepsilon(\theta|\Xi_\sigma) = P(d_\theta \leq \varepsilon|\theta), \tag{7}$$

where the approximation threshold $\varepsilon \in [0, \infty)$. We note that this approximation has the property that if $\delta(\Xi_\sigma^A, \Xi_\sigma^B) = 0 \Rightarrow \Xi_\sigma^A = \Xi_\sigma^B$, then at the limit of $\varepsilon = 0$, it matches the exact likelihood. This follows from

$$\begin{aligned}
\tilde{L}_0(\theta|\Xi_\sigma) &= P(d_\theta = 0|\theta) && (8)\\
&= P(\texttt{RLSUM}(M_\theta, |\Xi_\sigma|, \sigma) = \Xi_\sigma|\theta) && (9)\\
&= P(\Xi_\sigma|\theta) && (10)\\
&= L(\theta|\Xi_\sigma), && (11)
\end{aligned}$$

where the step from Equation 9 to Equation 10 is justified as $\texttt{RLSUM}(M_\theta, |\Xi_\sigma|, \sigma)$ is also the assumed generative process for Ξ_σ.

We note that a benefit of this approach is that we do not have to define the transfer function of the MDP as a probability distribution $P(s'|s, a)$; the only requirement is that we are able to generate $s' \sim P(s'|s, a)$. We also do not need to be able to evaluate $P\big(\sigma(\xi_i) = \xi_{i\sigma}\big)$ or assume that A or S are finite in size.

Algorithm 2 Exact Maximum Likelihood Inference Algorithm for IRL-SD

Input: M, Ξ_σ, Θ, H, N_{opt}
Output: $\hat{\theta}_{ML}$
$D \leftarrow \varnothing$
for $i = 1$ **to** N_{opt} **do**
  $\theta_i \leftarrow \arg\min_\theta Acq(\theta|D, H)$
  $\pi^*_{\theta_i} \leftarrow \texttt{RL}(M_{\theta_i})$
  $l_\theta \leftarrow -\log L(\theta_i|\Xi_\sigma)$
  $D \leftarrow \{D, (\theta_i, l_\theta)\}$
end for
$\hat{\theta}_{ML} \leftarrow \arg\min_\theta G_\mu(\theta|D, H)$

Algorithm 3 Approximate Maximum Likelihood Inference Algorithm for IRL-SD

Input: M, Ξ_σ, Θ, H, N_{opt}
Output: $\hat{\theta}_{ML}$
$D \leftarrow \varnothing$
for $i = 1$ **to** N_{opt} **do**
  $\theta_i \leftarrow \arg\min_\theta Acq(\theta|D, H)$
  $\Xi_\sigma^{sim} \leftarrow \texttt{RLSUM}(M_{\theta_i})$
  $d_\theta \leftarrow \delta(\Xi_\sigma^{sim}, \Xi_\sigma)$
  $D \leftarrow \{D, (\theta_i, d_\theta)\}$
end for
$\hat{\theta}_{ML} \leftarrow \arg\min_\theta G_\mu(\theta|D, H)$

4.3. Inference

Normal strategy for gradient-based methods would be to derive the gradient of the log-likelihood. However, in our case this is not feasible, as the likelihood has a sum-term and the surrogate contains arbitrary subroutines. For exact likelihood, MCMC could be possible in principle, but the high sample complexity of MCMC may make the computation infeasible in practice.

As evaluating even the surrogate is expensive, we use Bayesian optimization (BO) (see Brochu et al., 2010) combined with a Gaussian process (GP) regression model (Rasmussen, 2004) for finding the maximum likelihood (ML) point estimate. For a GP fit with data D and hyperparameters H, we denote the mean at θ by $G_\mu(\theta|D, H)$ and the standard deviation by $G_\sigma(\theta|D, H)$. Regarding BO, we denote the optimization horizon by N_{opt} and the acquisition function value at θ by $Acq(\theta|D, H)$.

ML inference based on the exact likelihood is illustrated in Algorithm 2, and ML inference based on the surrogate likelihood is illustrated in Algorithm 3. In the approximate method we use the discrepancy values instead of surrogate likelihood values, as with mild assumptions they have the same optima but d_θ is numerically better behaved and does not require ε to be defined for evaluation.

For performing MAP or full posterior inference, one needs to further assume a prior $P(\theta)$. For exact posterior inference, one could replace the log-likelihood in Algorithm 2 with log-posterior and return $G_\mu(\theta|D,H)$ instead of $\hat{\theta}_{ML}$. In the case of approximate inference it is difficult to sample $\tilde{L}_\varepsilon$ directly. Here we take a similar approach as Gutmann & Corander (2016) by first fitting the GP to d_θ, then selecting a suitable value for ε (we used $\min_\theta G_\mu(\theta|D,H)$) and returning $\tilde{P}(\theta|\Xi_\sigma) = P(\theta)\tilde{L}_\varepsilon(\theta|\Xi_\sigma)$ with the distribution of d_θ estimated from the GP.

4.4. Computational Complexity

To simplify analysis, we assume a computational budget has been decided for various operations.[6] We denote the budget for solving the RL problem[7] by O_{RL}, and the budget for minimizing a bounded non-convex function f by O^f_{min}. We also assume that the costs of evaluating $P\big(\sigma(\xi_i) = \xi_{i\sigma}\big)$ and $\sigma(\xi)$ are constant.

Complexity of fitting the GP is $O(N^3_{opt})$ and the complexity of evaluating the mean and variance of the GP after fitting is $O(N^2_{opt})$ (Rasmussen, 2004). We assume here that these constitute the complexity for evaluating $Acq(\theta|G,H)$.[8]

Thus, the total complexity for performing exact inference is $O(N^4_{opt} + N^3_{opt}O^{Acq}_{min} + N_{opt}(O_{RL} + NO_{ap} + N|\Xi'_{ap}|T_{max})))$, where O_{ap} is the worst-case complexity of finding Ξ_{ap} given a summary observation and $|\Xi'_{ap}|$ is the size of the largest possible set Ξ_{ap} given a summary observation. Similarly, the total complexity for performing approximate inference is $O(N^4_{opt} + N^3_{opt}O^{Acq}_{min} + N_{opt}(O_{RL} + N)))$.

It is challenging to bound N_{opt} without any knowledge of the shape and dimensionality of the likelihood function. Recently Kawaguchi et al. (2015) demonstrated that regret can be made to converge exponentially as a function of N_{opt}, although the base number may be very close to 1.

Broadly speaking, given a constant computational budget, with both methods there is a trade-off between the size of the MDP (increases O_{RL} and $|\Xi'_{ap}|$) and number of parameters that can be estimated (increases N_{opt}).

5. Experiments

To study the performance of the proposed algorithms, we start with a well-known toy problem, where we have an agent navigating on a grid, but change the observation assumptions to match the IRL-SD problem. Through this example, we demonstrate that we are able to infer the parameters of the agent's reward function based on only observing the initial state and total number of steps of the paths the agent chose to reach the goal. We also demonstrate that the approximate method is able to achieve quality of results comparable to the exact method, but considerably faster.

After this, we demonstrate that our approach is applicable for realistic modeling cases as well. We show that the approximate method is able to infer reasonable parameter values for a RL-based user model from HCI literature, based on measurements of real user behavior. Additional technical details of the experiments are collected in Section 5.5.

5.1. Grid World

Grid world is a well-known problem type in IRL literature (Ng & Russell, 2000; Abbeel & Ng, 2004; Neu & Szepesvári, 2007; Boularias et al., 2011; Herman et al., 2016). In this problem, an agent is located on a cell in a discrete two-dimensional grid. The agent has four actions that allow it to move to neighboring cells. Each action fails with probability p_{slip}, resulting in the agent moving to a random neighboring cell. When the agent enters a cell, it receives a reward based on the features of the cell $\phi(s)$ and the features of the agent's reward function θ, according to $R(s) = \theta^T\phi(s) + r_{step}$. Attempting to move outside the grid returns the agent to the current cell.

In our variant, we use a square grid of $w \times w$ cells, with $p_{slip} = 0.05$ and $r_{step} = -0.05$. The agent is initially located on a random cell at the edge of the grid. The cell at the center of the grid is the goal, and entering the goal gives the agent reward 1.0 and ends the episode. T_{max} was set to $3w$ steps, which was assumed to be sufficient for nearly any reasonable strategy.

Each grid cell has N_f binary features. The features of the grid cells are generated by placing w walls for each feature at random on the grid.[9] A wall is placed by (1) choosing one grid cell at random; (2) choosing vertical or horizontal direction at random; (3) choosing another cell at random along the chosen direction; (4) setting the feature values to 1 for all cells between (and including) these two cells, except if the cell is the goal cell. An example of a grid with three features generated this way is shown in Figure 1.

The summary function is defined as $\sigma(\xi) = (s_0, |\xi|)$, yielding the initial state, and the number of steps it took for the agent to reach the goal (i.e. we do not know what the intermediate states or actions were). Our problem is to infer likely values for $\theta \in [-1, 0]^{N_f}$ given a set of summary observations Ξ_σ and the MDP definition.

[6] Feasible budgets might be estimated based on initial tests.

[7] Regret can be bounded by $\tilde{O}(D|S|\sqrt{|A|T})$, where D is the diameter of the MDP (Jaksch et al., 2010).

[8] This is exact if we use the LCB acquisition function.

[9] The seed value used for generating the grid is part of the MDP definition.

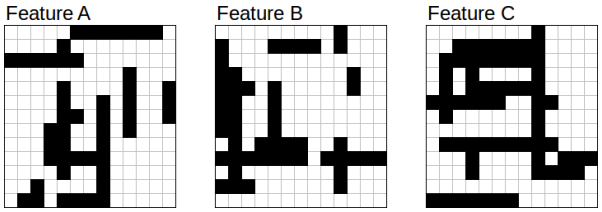

Figure 1. Visualization of a 13×13 grid with three features, generated by placing 13 random walls per feature. Each feature is shown individually, with black squares denoting presence of the feature. Each feature can be thought of as a different type of obstacle (e.g. mountains, swamp, forest).

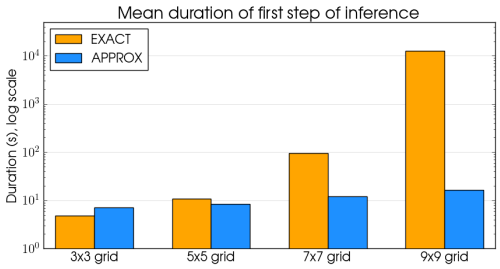

Figure 2. Time it takes to run the first step of the exact and approximate ML estimation algorithms (in logarithmic scale). Mean of 20 independent experiments.

5.1.1. Exact inference

As the number of possible paths given an observation is finite, the likelihood function can be evaluated in principle. Given a start state and path length, the paths form a tree with branching factor at most 16 and depth equal to the path length.[10] The tree can be pruned by requiring that the absorbing goal state appears exactly once in the path, as the final state. Also, as in this case we use a deterministic policy, we can prune out all actions for each state that disagree with the policy.

5.2. Experiment 1: Estimating Runtime

In this experiment, we compare the empirical runtimes of the exact and approximate method. We simulated observation sets with $N = 1000$ from grids of various sizes. We used grids with no features ($N_f = 0$) to avoid long paths that would make the exact method infeasible to evaluate. We computed the first iteration step for both algorithms and recorded the elapsed wall-clock time. The results are shown in Figure 2.

We observe that the empirical runtime of the exact algorithm grows rapidly as the size of the grid increases. This is expected, as $|\Xi_{ap}|$ grows exponentially as the length of the path grows linearly. On the other hand, the runtime of the approximate algorithm is comparatively nearly constant, only increasing slightly because the RL step takes longer as the grid size grows.

5.3. Experiment 2: Inference Quality

In this experiment we compare the quality of ML inference between the exact and approximate method on small grids. We also investigate the performance of the approximate method on larger grids, where the exact method would be infeasible to compute.

[10] There are at most 4 optimal actions at each state and at most 4 possible succeeding states.

The inference was done based on observation sets with $N = 1000$. We performed inference with N_f being 2 and 3. The number of BO samples was $200N_f$, computed in batches of $20N_f$. The locations of the first 20 % of the samples were chosen uniformly at random. The discrepancy between simulated and observation datasets was measured with the mean absolute error in path length per individual starting location. We also measured the ML inference error with L_2-distance to the ground truth.

To bound the computation time of the exact method, we excluded any observations from Ξ_σ where the path contained more than 10 actions (7×7 grid) or 13 actions (9×9 grid), leaving on average 77 % of the observations. The same Ξ_σ was used for both algorithms when comparing performance between the exact and approximate method.

We observe that the approximate method compares favorably to the exact method, shown as the smaller discrepancy and L_2 error achieved by the approximate method on 7×7 and 9×9 grids in Figure 3. It is understandable that the approximate method achieves lower discrepancy, as it has been formulated to minimize this quantity. However, it is surprising that the approximate method also achieves lower L_2 error to the ground truth. One reason for this is that the precise optimum of the log-likelihood is more difficult to find, as the magnitude of variation near the optimum is relatively smaller with the log-likelihood values compared to the discrepancy values.

We also observe that the approximate method is able to perform well even on larger grids where the exact method would be infeasible to compute. In Figure 3 the discrepancy of the predictions increases predictably with the grid size, and the L_2 error is comparable to the performance with smaller grids. The runtime of the approximate method seems to also grow at a predictable rate when the size of the grid increases, shown in Figure 4.

In general, when measured with L_2 error, both methods are able to produce inference results that are clearly bet-

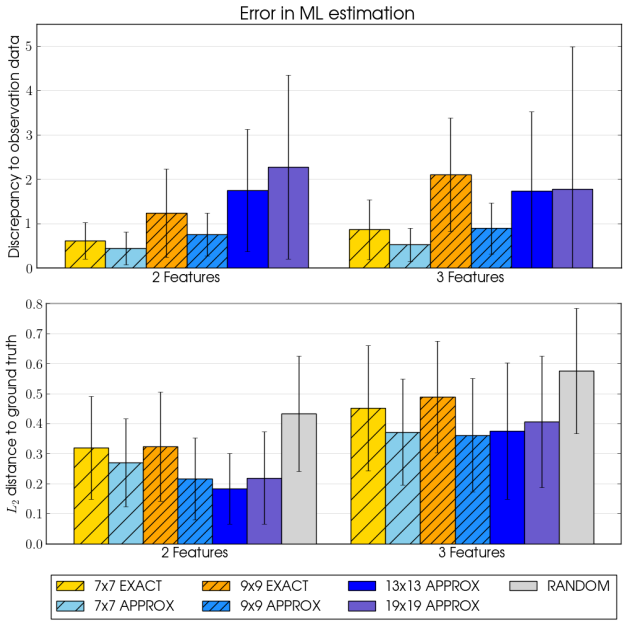

Figure 3. Quality of ML estimates with the exact and approximate inference methods. Top: discrepancy to the observation data, smaller is better. Estimated by selecting the sample from the approximate method experiment that was closest to the final ML location. Bottom: L_2 distance to ground truth, smaller is better. Random baseline represents a uniform guess from the parameter space. Mean and standard deviation of 30 independent experiments.

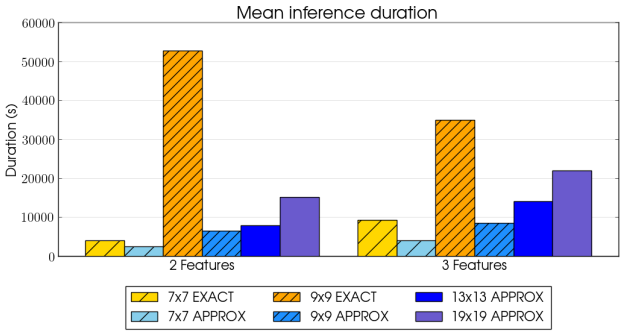

Figure 4. The duration of inference with the exact and approximate inference methods, measured in wall-clock time. N_{opt} was 400 when inferring 2 features and 600 when inferring 3 features. Mean and standard deviation of 30 independent experiments. There was significant variation in the duration of the inference for the 9×9 grid with the exact method. 10k seconds is approximately 3 hours.

ter than a random baseline. Also, both methods are able to produce ML estimates that lead to reasonably small discrepancy measures with the observation data (average difference in path length being at most 2 steps per starting location).

5.4. Experiment 3: Modelling Computer Users

In the final experiment we study the performance of the approximate method on a realistic model, for which we infer the full posterior based on real observation data.

The task is to estimate the parameters of a MDP where the goal of the user is to search for an item from a computer drop-down menu (Chen et al., 2015; Kangasrääsiö et al., 2017). Although with large screens traditional IRL methods have been used (Mohammed & Staadt, 2015) as detailed actions can be measured with eye-tracking, with small menus the accuracy of eye-tracking is often poor in comparison. However, simple summary statistics, such as the time between opening a menu and clicking the target item, are simple to measure accurately, but require solving the IRL-SD problem.

5.4.1. Model Definition

In this model the drop-down menu consists of 8 items, which can be in multiple states: the relevance of the item is either unobserved, low, medium or high, or the item is the target item. The length of the item either unobserved, correct or incorrect. In addition to this, the eyes of the agent are fixated on one of the items. These constitute the state that the agent observes.

Initially all items are unobserved, and the agent can observe them by fixating at any item. With some probability the agent will observe the length or semantic relevance of the neighboring items as well during a fixation. Also, with probability p_{rec} the agent will recall the menu layout after the first fixation.[11] The duration of a fixation is f_{dur} (in units of 100 ms). The agent can select the item that is being observed, which takes d_{sel} seconds. The agent can also quit, which is instantaneous.

The menu either contains or does not contain the target item. If the menu happens to contain the target, the agent receives a large reward when she finds the target item and selects it. If the menu does not contain the target item, the agent receives a large reward when she quits. Taking the wrong end action results in a large penalty. After each action, the agent incurs a penalty according to the time spent.

5.4.2. Inference

Recently Kangasrääsiö et al. (2017) estimated the MAP values of the parameters of this model using summary observations from a user study by Bailly et al. (2014). The

[11] That is, observe the features of all the items.

summary observation included the task completion time (TCT, sum of the durations of all actions in an episode) and whether the target was present or absent in the menu. Here we extend their analysis by showing that full posteriors can be estimated based on the same dataset and model.

We use the same discrepancy function as in the original paper, based on the differences in TCT means and standard deviations in each menu condition.

We infer the posteriors of three parameters: f_{dur}, d_{sel} and p_{rec}. The rest of the model parameters are fixed to the MAP values inferred by Kangasrääsiö et al. (2017). We use the truncated normal distribution as prior for two of the parameters: similarly as in the original paper, for f_{dur} we use mean 3, std 1; for d_{sel} mean 0.3, std 0.3. For p_{rec} we use a uniform prior. To make visualization easier, we infer the posterior in two groups: first jointly for f_{dur} and d_{sel} (Figure 5 left), then jointly for f_{dur} and p_{rec} (Figure 5 right). We computed 400 BO samples in batches of 80 for estimating each of the posteriors. The locations of the first 80 samples were chosen uniformly at random. The inference took 3 days for each posterior, reinforcement learning taking roughly 12 h per sample.

5.4.3. Results

We observe that *a posteriori* there is clear correlation between f_{dur} and d_{sel}, and similarly for f_{dur} and p_{rec}. Both of these are understandable, as increasing f_{dur} would increase the predicted TCT, as would decreasing p_{rec} or increasing d_{sel}. The posterior of f_{dur} seems to be quite tightly centered around 200 ms, but there is considerable uncertainly left in d_{sel} and in p_{rec}. The uncertainty in d_{sel} could be explained by the difficulty of pointing precisely to the target item with the cursor, which may cause variation in its duration. The uncertainty in p_{rec} could be explained by the fact that the menus encountered early on in the experiments were completely new to the subjects, but at the end of the experiment the subjects were more likely to recall a previously encountered menu. We note that these insights would not have been possible to infer from just MAP estimates.

5.5. Technical Details

Here we have collected the technical details related to the experiments. In the grid world experiments, the ground truth values for θ were fixed to be evenly spread out over the range of possible values: $[-0.33, -0.67]$ (2 features) or $[-0.25, -0.5, -0.75]$ (3 features). This was done to improve the identifiability of the parameters and reduce the amount of noise caused by variation in the experiment setting. Optimal policy was learned with Q-learning over 10k episodes in Experiment 1 and 100k episodes in Experiment 2. In the menu model, the optimal policy was learned with Q-learning over 20M training sessions. We used a fixed set of 10k training menus during the training phase, and a fresh set of 10k menus for the simulation phase.

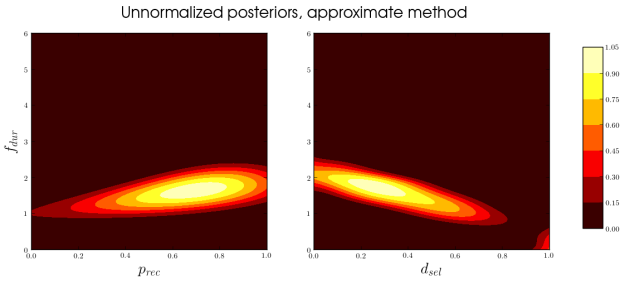

Figure 5. Left: posterior of fixation duration f_{dur} and selection delay d_{sel}. Right: posterior of fixation duration f_{dur} and menu recall probability p_{rec}.

A radial basis function kernel was used for the GP. The initial lengthscale of the kernel was set to 20 % of the smallest bound width and the variance to roughly 50 % of the maximum difference between the minimum and maximum sample values observed in initial tests. After each batch, the lengthscale and variance were optimized from the initial values, based on the current observation data. The noise variance was set to 1 % of the kernel variance for the exact method and to 10 % for the approximate method (as evaluations are much more noisy). We used the same type of acquisition function as Gutmann & Corander (2016): we first found $\theta_i = \arg\min_\theta Acq(\theta|D, H) = G_\mu(\theta|D, H) - G_\sigma(\theta|D, H)$ and then estimated the curvature of Acq at θ_i. The locations for the samples in the batch were drawn from a normal distribution centered at θ_i with the same curvature at the maximum.

6. Discussion

In this paper we defined the IRL-SD problem, where the task is to do inverse reinforcement learning based on summary observations of the agent's behavior. We proposed exact and approximate methods for inference in this situation. We demonstrated that the exact method is computationally very expensive, but the approximate method can be used even for full posterior inference with complex MDPs and real observation data.

The selection of the discrepancy function δ and threshold ε is generally an open question. A recent summary of different methods for choosing the discrepancy function and the threshold are provided in (Lintusaari et al., 2017). There exists also recent work where a classifier is used as the discrepancy function (Gutmann et al., 2017), which is a general approach that might work with several models.

Acknowledgements

This work has been supported by the Academy of Finland (Finnish Centre of Excellence in Computational Inference Research COIN, and grants 294238, 292334). Computational resources were provided by the Aalto Science IT project.

References

Abbeel, Pieter and Ng, Andrew Y. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the 21st International Conference on Machine Learning*, pp. 1–8, 2004. doi: 10.1145/1015330.1015430.

Abbeel, Pieter, Coates, Adam, and Ng, Andrew Y. Autonomous helicopter aerobatics through apprenticeship learning. *The International Journal of Robotics Research*, 29(13):1608–1639, 2010. doi: 10.1177/0278364910371999.

Bailly, Gilles, Oulasvirta, Antti, Brumby, Duncan P., and Howes, Andrew. Model of visual search and selection time in linear menus. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, pp. 3865–3874, 2014. doi: 10.1145/2556288.2557093.

Banovic, Nikola, Buzali, Tofi, Chevalier, Fanny, Mankoff, Jennifer, and Dey, Anind K. Modeling and understanding human routine behavior. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pp. 248–260, 2016. doi: 10.1145/2858036.2858557.

Bloem, Michael and Bambos, Nicholas. Infinite time horizon maximum causal entropy inverse reinforcement learning. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pp. 4911–4916, 2014. doi: 10.1109/CDC.2014.7040156.

Boularias, Abdeslam, Kober, Jens, and Peters, Jan. Relative entropy inverse reinforcement learning. In *JMLR Workshop and Conference Proceedings Volume 15: AISTATS 2011*, pp. 182–189, 2011.

Brochu, Eric, Cora, Vlad M., and De Freitas, Nando. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. Technical Report TR-2009-023 and arXiv:1012.2599, University of British Columbia, 2010.

Chandramohan, Senthilkumar, Geist, Matthieu, Lefevre, Fabrice, and Pietquin, Olivier. User simulation in dialogue systems using inverse reinforcement learning. In *Interspeech*, pp. 1025–1028, 2011.

Chen, Xiuli, Bailly, Gilles, Brumby, Duncan P, Oulasvirta, Antti, and Howes, Andrew. The emergence of interactive behavior: A model of rational menu search. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 4217–4226, 2015. doi: 10.1145/2702123.2702483.

Choi, Jaedeug and Kim, Kee-Eung. Map inference for Bayesian inverse reinforcement learning. In *Advances in Neural Information Processing Systems 24*, pp. 1989–1997, 2011.

Choi, Jaedeug and Kim, Kee-Eung. Hierarchical Bayesian inverse reinforcement learning. *IEEE Transactions on Cybernetics*, 45(4):793–805, 2015. doi: 10.1109/TCYB.2014.2336867.

Dimitrakakis, Christos and Rothkopf, Constantin A. Bayesian multitask inverse reinforcement learning. In *Proceedings of the 9th European Workshop on Recent Advances in Reinforcement Learning*, pp. 273–284, 2011. doi: 10.1007/978-3-642-29946-9_27.

Gao, Yang, Peters, Jan, Tsourdos, Antonios, Zhifei, Shao, and Meng Joo, Er. A survey of inverse reinforcement learning techniques. *International Journal of Intelligent Computing and Cybernetics*, 5(3):293–311, 2012. doi: 10.1108/17563781211255862.

Gutmann, Michael U. and Corander, Jukka. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research*, 17(1):4256–4302, 2016.

Gutmann, Michael U., Dutta, Ritabrata, Kaski, Samuel, and Corander, Jukka. Statistical inference of intractable generative models via classification. *Statistics and Computing*, 2017. In press, preprint arXiv:1407.4981.

Herman, Michael, Gindele, Tobias, Wagner, Jörg, Schmitt, Felix, and Burgard, Wolfram. Inverse reinforcement learning with simultaneous estimation of rewards and dynamics. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pp. 102–110, 2016.

Jaksch, Thomas, Ortner, Ronald, and Auer, Peter. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(Apr):1563–1600, 2010.

Kangasrääsiö, Antti, Athukorala, Kumaripaba, Howes, Andrew, Corander, Jukka, Kaski, Samuel, and Oulasvirta, Antti. Inferring cognitive models from data using approximate Bayesian computation. In *Proceedings of the 35th Annual ACM Conference on Human Factors in Computing Systems*, 2017. doi: 10.1145/3025453.3025576. In press, preprint arXiv:1612.00653.

Kawaguchi, Kenji, Kaelbling, Leslie Pack, and Lozano-Pérez, Tomás. Bayesian optimization with exponential convergence. In *Advances in Neural Information Processing Systems 28*, pp. 2809–2817, 2015.

Kitani, Kris M., Ziebart, Brian D., Bagnell, James Andrew, and Hebert, Martial. Activity forecasting. In *Proceedings of the 12th European Conference on Computer Vision*, pp. 201–214, 2012. doi: 10.1007/978-3-642-33765-9_15.

Klein, Edouard, Geist, Matthieu, Piot, Bilal, and Pietquin, Olivier. Inverse reinforcement learning through structured classification. In *Advances in Neural Information Processing Systems 25*, pp. 1007–1015, 2012.

Klein, Edouard, Piot, Bilal, Geist, Matthieu, and Pietquin, Olivier. A cascaded supervised learning approach to inverse reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 1–16, 2013. doi: 10.1007/978-3-642-40988-2_1.

Lintusaari, Jarno, Gutmann, Michael U., Dutta, Ritabrata, Kaski, Samuel, and Corander, Jukka. Fundamentals and recent developments in approximate Bayesian computation. *Systematic Biology*, 66(1):e66, 2017. doi: 10.1093/sysbio/syw077.

Michini, Bernard and How, Jonathan P. Bayesian nonparametric inverse reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 148–163, 2012. doi: 10.1007/978-3-642-33486-3_10.

Mohammed, Redwan Abdo A. and Staadt, Oliver. Learning eye movements strategies on tiled large high-resolution displays using inverse reinforcement learning. In *International Joint Conference on Neural Networks*, pp. 1–7, 2015. doi: 10.1109/IJCNN.2015.7280675.

Neu, Gergely and Szepesvári, Csaba. Apprenticeship learning using inverse reinforcement learning and gradient methods. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pp. 295–302, 2007.

Ng, Andrew Y. and Russell, Stuart. Algorithms for inverse reinforcement learning. In *Proceedings of the 17th International Conference on Machine Learning*, pp. 663–670, 2000.

Nguyen, Quoc Phong, Low, Bryan Kian Hsiang, and Jaillet, Patrick. Inverse reinforcement learning with locally consistent reward functions. In *Advances in Neural Information Processing Systems 28*, pp. 1747–1755, 2015.

Ramachandran, Deepak and Amir, Eyal. Bayesian inverse reinforcement learning. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, volume 51, pp. 1–4, 2007.

Rasmussen, Carl Edward. Gaussian processes in machine learning. In *Advanced Lectures on Machine Learning*, pp. 63–71. 2004. doi: 10.1007/978-3-540-28650-9_4.

Ratliff, Nathan D., Bagnell, J. Andrew, and Zinkevich, Martin A. Maximum margin planning. In *Proceedings of the 23rd International Conference on Machine Learning*, pp. 729–736, 2006. doi: 10.1145/1143844.1143936.

Rothkopf, Constantin A. and Dimitrakakis, Christos. Preference elicitation and inverse reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 34–48, 2011. doi: 10.1007/978-3-642-23808-6_3.

Russell, Stuart. Learning agents for uncertain environments. In *Proceedings of the 11th Annual Conference on Computational Learning Theory*, pp. 101–103, 1998. doi: 10.1145/279943.279964.

Sunnåker, Mikael, Busetto, Alberto Giovanni, Numminen, Elina, Corander, Jukka, Foll, Matthieu, and Dessimoz, Christophe. Approximate Bayesian computation. *PLOS Computational Biology*, 9(1):1–10, 01 2013. doi: 10.1371/journal.pcbi.1002803.

Surana, Amit. Unsupervised inverse reinforcement learning with noisy data. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pp. 4938–4945, 2014. doi: 10.1109/CDC.2014.7040160.

Tossou, Aristide C. Y. and Dimitrakakis, Christos. Probabilistic inverse reinforcement learning in unknown environments. In *Conference on Uncertainty in Artificial Intelligence*, 2013.

Wobbrock, Jacob O., Kane, Shaun K., Gajos, Krzysztof Z., Harada, Susumu, and Froehlich, Jon. Ability-based design: Concept, principles and examples. *ACM Transactions on Accessible Computing*, 3(3):9, 2011. doi: 10.1145/1952383.1952384.

Ziebart, Brian D., Maas, Andrew L., Bagnell, J. Andrew, and Dey, Anind K. Maximum entropy inverse reinforcement learning. In *Proceedings of the National Conference on Artificial Intelligence*, pp. 1433–1438, 2008.

Ziebart, Brian D., Ratliff, Nathan, Gallagher, Garratt, Mertz, Christoph, Peterson, Kevin, Bagnell, J. Andrew, Hebert, Martial, Dey, Anind K., and Srinivasa, Siddhartha. Planning-based prediction for pedestrians. In *International Conference on Intelligent Robots and Systems*, pp. 3931–3936, 2009. doi: 10.1109/IROS.2009.5354147.